

123456789101112131415161718192021222324252627282930313233343536373839404142434445464748495051525354555657585960616263646566676869707172737475767778798081828384858687888990919293949596979899100

1. Field of the Invention

5

10

15

20

25

only those receivers which are part of that multicast group. Many IP Multicast protocols exist which differ essentially in the way the distribution trees are created and shared.

The Distance Vector Multicast Routing Protocol (DVMRP), one of the oldest multicast routing protocols, belongs to the category of broadcast-and-prune protocols where the multicast data itself is used to build the tree. The DVMRP protocol assumes that all routers in a network want to receive all multicast packets. Packets are usually flooded to all routers, and those routers that do not want to receive multicast transmissions belonging to a particular group must send an explicit Prune message to their upstream routers to indicate this. DVMRP creates a (source, group) pair at every router. A shortest-path tree is created for every source of every group, to connect that source with all the other members of the group.

Protocol Independent Multicast-Sparse Mode (PIM-SM) is another
15 protocol which is being increasingly deployed in the Internet. PIM-SM builds
shared trees which must be explicitly joined by downstream routers. All the
sources in a group share the tree which is built around a central routing entity;
i.e., a "core" or "Rendezvous Point (RP)." Every router that wants to join a
particular multicast group must send an explicit Join to the RP for that group.
20 This protocol is in contrast to broadcast-and-prune protocols discussed above,
which broadcast multicast traffic downstream until explicit prunes are
received. Further details of IP multicast, DVMRP, and PIM-SM can be found
in publications such as: T. Maufer, "Deploying IP Multicast in the Enterprise,"
Prentice Hall, 1998; D. Waitzman, C. Partridge, and S. Deering, "Distance
25 Vector Multicast Routing Protocol," RFC 1075, November 1988; and Estrin,
D., Farinacci, D., Helmy, A., Thaler, D., Deering, S., Handley, M., Jacobson,
V., Liu, C., Sharma, P., and L. Wei, "Protocol Independent Multicast-Sparse
Mode (PIM-SM): Protocol Specification," RFC 2117, June 1997.

In mesh-TDMA satellite networks, the communication is usually point-to-point, and each transmit-receive pair of terminals is assigned one or

more slots (bursts) in a frame for communication. Channel capacity is increased by adding carriers, and, in advanced TDMA schemes, terminals can hop carriers from burst-to-burst so as to maximize utilization of the space segment. Efficient multicasting in such an environment requires the
5 identification of a slot/burst in which the closest terminal to a multicast source can transmit, and during which all receive terminals which are part of the group can listen. Finding such available slots can become a formidable problem which only worsens if group membership is dynamic.

The most straightforward technique for supporting multicast IP routing
10 and forwarding in a mesh satellite network would be to incorporate multicast routing/forwarding capabilities into each terminal. A packet-switched satellite terminal typically has one or more terrestrial interfaces (such as X.25, frame-relay, ATM, or ethernet) and a single physical satellite interface. The satellite interface can be used to communicate with one, many, or all of the terminals
15 in a network depending on the beam connectivity and available bandwidth on the satellite. Since multicast routing messages are typically exchanged between a router and all of its adjacent neighbors, the terminal/router would need to periodically communicate with all of the terminal/routers in the mesh, thereby using satellite link bandwidth. Supporting multicast routing protocols
20 at the terminal can require significant CPU and memory resources. In the case of DVMRP, prune-state tables need to be maintained in all the terminals, which could potentially consume a lot of memory. In addition, significant effort may be required to port and test multicast routing protocol software, especially if multiple protocols were to be supported.

25 IP multicast has been reportedly deployed in hub-spoke type satellite networks, wherein satellite terminals communicate with a hub and not with each other. The hub broadcasts all its outbound data on a single carrier which is received by all terminals. Terminals demultiplex data intended for them and discard the rest. However, IP multicast has not been publicly deployed on
30 mesh satellite networks.

SUMMARY OF THE INVENTION

It is, therefore, an object of the present invention to provide a bandwidth-efficient technique for routing multicast IP traffic over meshed satellite networks.

5 This and other objects are achieved by a system architecture wherein the multicast routing protocols are run in a centralized route server (implemented on a standard UNIX workstation). In this architecture, the satellite network is part of a router fabric, with terminals appearing as ports attached to the router core (the route-server). In the baseline implementation, 10 external hosts/subnets connect to terminals through multicast enabled routers. External routers establish multicast routing sessions only with the route-server, and not with other terminals. Multicast routing packets originated by an external router attached to a terminal will be conveyed transparently to the route-server and used to create multicast group table information at the route 15 server. This information is provided by the route -server to the terminals so that multicast traffic can be directly transmitted from the ingress terminal to all the terminals in a group, without having to be relayed through the route server.

Implementing the route-server on a workstation provides enough CPU and memory resources to run common routing protocols and store large 20 routing tables. The route-server can be easily upgraded with more memory and extra processing power if necessary.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention will be more clearly understood from the following description in conjunction with the accompanying drawing, wherein:

25 Fig. 1 is a diagram of a mesh satellite network supporting multicast IP services in accordance with the present invention.

Fig. 2 is a flow chart illustrating the basic steps necessary to implement the invention.

Fig. 3 is a representation of slot assignment in a prior art TDMA system.

Fig. 4 is a representation of slot assignment according to the present invention.

5 Fig. 5 is a flow chart illustrating the additional steps necessary to implement the present invention using the DVMRP routing protocol.

Fig. 6 is a flow chart illustrating the additional steps necessary to implement the present invention using the PIM-SM routing protocol.

DETAILED DESCRIPTION OF THE INVENTION

10 An end-to-end network configuration which supports IP services over a mesh satellite network is shown in Figure 1.

Such a network can be used to provide connectivity via satellite-1 12 and satellite dishes 14, 16, 18 to Internet Service Providers, e.g., 10 or 20, and can also be used to connect corporate sites. Routers 52, 54, 56 and 58 connect
15 the terminals 34 and 36, possibly via frame relays 60, 62, 64 (or via an ethernet connection), to remote access servers 70, 72, and/or to various dial-up access units 101-110. The network control center (NCC) 30 is located at one site and it is typically a workstation which runs software responsible for configuring, controlling and monitoring the entire network of satellite
20 terminals. The terminal 32 at that site is like any other terminal 34, 36 in the network, but is referred to as the master terminal for clarity. Such a network configuration is typical of most mesh satellite networks. An addition to the normal network configuration is the route-server (RS) 40, which can be connected by router-1 50 via a WAN link to the Internet 68. The RS computer
25 is on the same Local Area Network (LAN) as the NCC 30 and the master terminal 32. It is assumed here that the beam connectivity is such that all terminals can communicate with each other directly (with a single satellite hop). However, straightforward extensions of the invention, such as the use of

two route servers (one in each area), could handle the case of directional beam connectivity.

One approach to the problem of allocating point-to-multipoint (PTM) bursts is to set aside one or more slots in a TDMA frame (on one or more carriers) to be only used for broadcast communications. Terminals which need to transmit data in a PTM burst will request allocation of broadcast bursts from the NCC 30, and all terminals in the network will listen to the broadcast bursts. In case there are very few receiving terminals for a particular multicast, and if broadcast capacity is to be conserved, specific PTM slots could be found by a burst time plan (BTP) generation algorithm. In case broadcast slots are used up, and there is more demand, additional broadcast slots can be assigned by taking away capacity from point-to-point traffic. The minimum and maximum amounts of broadcast capacity to be made available would be pre-configurable by the network operator. If no broadcast or PTM capacity is available, then the terminal will replicate packets and transmit them over point-to-point bursts.

A flowchart demonstrating the operation of the above-described configuration is shown in Figure 2.

In step S201, a master routing table is established in the RS 40 by communication between the RS and other routers 52, 54, 56, 58. In step S202, a terminal 34 requests multicast services from the NCC 30. In step S203, the NCC 30 then assigns a TDMA slot for broadcast services for the use of the requesting terminal, and updates the master routing table accordingly. Then, in step S204, the multicast services are broadcast by the originating, or ingress, terminal 34 over the assigned slot, so that all terminals receive the transmitted message.

Figure 3 demonstrates a conceptualization of prior art TDMA slots for unicast and multicast services.

In slots 1-3, routing information is updated between the various routers so that slots 4-6 can each be assigned to communication between two points;

09689738 101300

e.g., between dial-up access units 101 to 104, 102 to 106 and 103 to 109. Such an arrangement is problematic because it requires an exchange of information between essentially all of the various routers, and this information must be updated periodically. This arrangement is even more problematic for
5 multicast services, which also require an extensive exchange of information between routers (slots 7-9), because multiple slots (10-12) are needed for the multicast services; i.e., 101 to 104, 101 to 106 and 101 to 109.

In contrast, as shown in Figure 4, the present invention allows for a reduction in the number of slots required for routing information updates (i.e.,
10 to slots 1 and 5). This reduction occurs due to the fact that routing information is exchanged only between each router and the RS 40, and not between all routers. In addition, the number of slots required for multicast services is drastically reduced (to only slot 6), because any one point in the network can broadcast to all other points over the broadcast time slot requested by a
15 terminal and assigned by the NCC 30.

It should be noted that Figures 3 and 4 are not intended as exact descriptions of time slot allocation in either the prior art or the present invention. That is, it is not to be understood from Figure 3 that exactly three time slots are required for routing information updates which will allow the
20 communication described in slots 4-6 therein. Similarly, it is not to be understood that only one time slot is required for the present invention to update routing information in the RS 40 and/or assign the broadcast slot 6. Such information will depend on the number of routers present, bandwidth availability, etc. Instead, Figures 3 and 4 are intended to conceptually
25 demonstrate that, by the use of a centralized route server and the assignment of a broadcast time slot for multicast services, link bandwidth can be effectively conserved, and delay can be significantly reduced.

Thus, in the proposed architecture, the RS 40 runs the multicast routing protocols, while the terminals 34, 36 actually short-cut the traffic to the
30 appropriate destinations. The flows of multicast routing and multicast traffic

creates a forwarding entry for the (*,G) pair and sends the Join message to the next upstream router (if necessary). An external router attached to a terminal will therefore unicast PIM_JOIN messages to the RS 40 in step S602 if its routing table indicates that the RP can be reached through the satellite network. The PIM-SM software in the RS 40 will process the PIM_JOIN message in the normal fashion and will send it to the next upstream router (either attached to a terminal) or on the terrestrial side. As in the DVMRP case, the RS 40 will send a forwarding table update message to the upstream terminal 36 in step S603 indicating that the terminal should use broadcast or PTM bursts for multicast packets directed towards that particular multicast group. The RS 40 will also send a message to the terminal 34 which originated the Join message to add a forwarding entry for the (*,G) pair. Prune messages and switches to source specific shortest path trees will be handled in a similar manner.

15 *Sending to a Group*

When a source transmits multicast packets to a group, the attached router first encapsulates the packets in PIM-SM-Register packets and unicasts them to the group's RP in step S604. Based on the join messages that the RP gets, packets are then multicast to the appropriate hosts in step S605. As discussed above, PIM_JOIN messages from leaf routers 52 have already triggered the creation of appropriate multicast forwarding entries in the ingress and egress satellite terminals, therefore multicast packets entering a terminal will be transmitted appropriately.

In conclusion, a route server based architecture such as that discussed above can provide support for many multicast routing protocols in an elegant fashion. The use of short-cut forwarding at terminals will minimize delay and conserve bandwidth. The burst allocation scheme will intelligently allocate broadcast and point-to-multipoint bursts to multicast traffic so as to conserve bandwidth.